



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2014-11

Do High-Frequency Financial Data Help Forecast Oil Prices? The MIDAS Touch at Work

by Christiane Baumeister, Pierre Guérin and Lutz Kilian

Bank of Canada Working Paper 2014-11

March 2014

Do High-Frequency Financial Data Help Forecast Oil Prices? The MIDAS Touch at Work

by

Christiane Baumeister,¹ Pierre Guérin¹ and Lutz Kilian²

¹International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9

²Department of Economics
University of Michigan
Ann Arbor, MI 48104
and
CEPR

Corresponding author: lkilian@umich.edu

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

We thank Frank Schorfheide for providing access to his MF-VAR code, Cynthia Wu for explaining the conventions used in constructing the weekly data used in Hamilton and Wu (2013), and Dmitri Tchebotarev for excellent research assistance.

Abstract

The substantial variation in the real price of oil since 2003 has renewed interest in the question of how to forecast monthly and quarterly oil prices. There also has been increased interest in the link between financial markets and oil markets, including the question of whether financial market information helps forecast the real price of oil in physical markets. An obvious advantage of financial and energy market data in forecasting oil prices is their availability in real time on a daily or weekly basis. We investigate whether mixed-frequency models can be used to take advantage of these rich data sets. We show that, among a range of alternative high-frequency predictors, changes in U.S. crude oil inventories produce substantial and statistically significant real-time improvements in forecast accuracy. The preferred mixed-data sampling (MIDAS) model reduces the mean-squared prediction error by as much as 16 percent compared with the no-change forecast and has statistically significant directional accuracy as high as 80 percent. This MIDAS forecast also is more accurate than a mixed-frequency real-time vector autoregressive forecast, but not systematically more accurate than the corresponding forecast based on monthly inventories. We conclude that typically not much is lost by ignoring high-frequency financial data in forecasting the monthly real price of oil.

JEL classification: C53, G14, Q43

Bank classification: Econometric and statistical methods; International topics

Résumé

La variation considérable des prix réels du pétrole depuis 2003 a ravivé l'intérêt porté aux méthodes de prévision des cours mensuels et trimestriels de ce produit. On a aussi observé un regain d'intérêt pour l'étude du lien entre les marchés financiers et pétroliers : à ce titre, les chercheurs se sont demandé si l'information en provenance des marchés financiers aide à prédire les prix réels du pétrole sur les marchés au comptant. Les données des marchés financiers et énergétiques présentent un avantage évident pour prévoir les cours du pétrole : elles sont accessibles en temps réel selon une fréquence quotidienne ou hebdomadaire. Nous cherchons donc à déterminer le pouvoir de prévision de ces riches ensembles de données en utilisant des modèles avec données à fréquence mixte. Nous montrons que, parmi toute une gamme de prédictors de haute fréquence, les variations des stocks de pétrole brut aux États-Unis améliorent de manière appréciable et statistiquement significative, en temps réel, l'exactitude des prévisions. Le modèle d'échantillonnage de données de fréquence mixte (MIDAS) privilégié peut réduire l'erreur quadratique moyenne de prévision dans une proportion allant jusqu'à 16 %, si l'on compare avec la prévision du modèle de marche aléatoire. Il permet également de prévoir avec exactitude le sens des variations dans 80 % des cas, et cela de manière statistiquement significative. Les prévisions établies grâce à ce modèle MIDAS sont également plus justes que celles issues d'un modèle vectoriel autorégressif basé sur des

données de fréquence mixte en temps réel, mais elles ne sont pas systématiquement plus exactes que les projections correspondantes fondées sur les stocks mensuels. Nous concluons que, généralement, le fait de ne pas tenir compte des données financières de haute fréquence dans la prévision des prix mensuels réels du pétrole a une incidence négligeable.

Classification JEL : C53, G14, Q43

Classification de la Banque : Méthodes économétriques et statistiques; Questions internationales

1 Introduction

The substantial variation in the real price of oil since 2003 has renewed interest in the question of how to forecast monthly and quarterly oil prices.¹ The links between financial markets and the price of oil have received particular attention, including the question of whether financial market information can help forecast the price of oil in physical markets (e.g., Fattouh, Kilian and Mahadeva 2013). An obvious advantage of financial data is their availability in real time on a daily basis. Financial data are not subject to revisions and are available on a daily or weekly basis. Existing forecasting models for the monthly real price of oil do not take advantage of these rich data sets. Our objective is to assess whether there is useful predictive information for the real price of oil in high-frequency data from financial and energy markets, and to identify which predictors are most useful. Incorporating daily or weekly financial data into monthly oil price forecasts requires the use of models for mixed-frequency data.

The development of models for variables sampled at different frequencies has attracted substantial interest in recent years. A comprehensive review can be found in Foroni, Ghysels and Marcellino (2013). A large and growing literature has documented the benefits of combining data of different frequencies in forecasting macroeconomic variables such as real GDP growth and inflation. One approach has been to construct mixed-frequency vector autoregressive (MF-VAR) forecasting models (e.g., Schorfheide and Song 2012). An alternative approach is the use of univariate mixed-data sampling (MIDAS) models (e.g., Andreou, Ghysels and Kourtellis 2011). The MIDAS model employs distributed lag polynomials to ensure a parsimonious model specification, while allowing for the use of data sampled at different frequencies. The original MIDAS model requires non-linear least-squares estimation (see Andreou, Ghysels and Kourtellis 2010). Foroni, Marcellino and Schuhmacher (2012) proposed a simplified version of the MIDAS model (referred to as unrestricted MIDAS or U-MIDAS) that can be estimated by ordinary least squares and in many applications has been shown to produce highly accurate out-of-sample forecasts,

¹A comprehensive review of this literature is provided in Alquist, Kilian and Vigfusson (2013). More recent contributions include Chen (2013), Baumeister and Kilian (2013a,b), Baumeister, Kilian and Zhou (2013), and Bernard, Khalaf, Kichian and Yelou (2013).

provided the data frequencies to be combined are not too different.

Numerous studies have documented the ability of MIDAS regressions to improve the accuracy of (1) quarterly macroeconomic forecasts based on monthly predictors and (2) monthly forecasts based on daily or weekly predictors (e.g., Andreou, Ghysels and Kourtellis 2013; Clements and Galvao 2008, 2009; Ghysels and Wright 2009; Hamilton 2008). Of particular interest are high-frequency financial data. One reason is that financial asset prices embody forward-looking information. Another reason is that financial data are accurately measured and available in real time, while lower-frequency macroeconomic data tend to be subject to revisions and become available only with a delay.

These differences in informational structure are particularly evident when forecasting oil prices. Commonly used predictors of the real price of oil, such as global oil production, global oil inventories, global real activity, or the U.S. refiners' acquisition cost for crude oil, only become available with considerable delays and are subject to potentially large but unpredictable revisions that can persist for up to two years (see Baumeister and Kilian 2012). Despite these drawbacks, several recent studies have shown that it is possible to systematically beat the no-change forecast of the monthly real price of oil in real time (e.g., Baumeister and Kilian 2012, 2013a,b).

The current paper investigates whether the accuracy of oil price forecasts can be improved by utilizing high-frequency information from financial markets and from U.S. energy markets. The set of high-frequency predictors includes (1) the spread between the spot prices of gasoline and crude oil; (2) the spread between the oil futures price and the spot price of crude oil; (3) cumulative percent changes in the Commodity Research Bureau (CRB) index of the price of industrial raw materials, (4) U.S. crude oil inventories, and (5) the Baltic Dry Index (BDI); (6) returns and excess returns on oil company stocks; (7) cumulative changes in U.S. nominal interest rates (LIBOR, federal funds rate), and (8) cumulative percent changes in the U.S. trade-weighted nominal exchange rate.

Our starting point is a MIDAS model for the monthly real price of oil. For reasons discussed in section 2, we focus on predictors measured at weekly intervals constructed from daily observations. As is standard in the oil price forecasting literature, we assess all forecasts based on their mean-squared prediction errors (MSPEs) and directional accuracy.

We consider forecast horizons, h , ranging from 1 month to 24 months. Our MIDAS models nest the no-change forecast of the real price of oil, allowing us to compare the accuracy of MIDAS regressions with that of competing models evaluated against the same benchmark. We also compare the MIDAS model forecasts to real-time forecasts from the corresponding model based on the same predictors measured at monthly frequency.

Our results reinforce and strengthen recent evidence that the monthly real price of oil is forecastable in real time. We find that the most accurate h -month-ahead forecasts are obtained by including the percent change in U.S. crude oil inventories over the preceding h months. The preferred MIDAS forecast has statistically significant directional accuracy as high as 71% at the 12-month horizon, for example, and as high as 77% at the 24-month horizon. It also produces statistically significant MSPE reductions relative to the no-change forecast of 8% at the 12-month horizon and 16% at the 24-month horizon. These improvements in forecast accuracy are very large by the standard of previous work on forecasting oil prices. At horizons below 12 months, however, the MSPE reductions of this MIDAS model are quite modest.

How the MIDAS model is implemented matters to some extent. While there is typically little difference in accuracy between the MIDAS model with equal weights and the MIDAS model with estimated weights, the unrestricted MIDAS model tends to be slightly less accurate than the other specifications. The success of these MIDAS forecasts based on U.S. crude oil inventories prompted us to also investigate the accuracy of the MF-VAR model obtained by including the same weekly inventory data in a monthly oil market VAR forecasting model of the type examined in Baumeister and Kilian (2012). We found that the latter specification did not perform systematically better than the original VAR model and clearly worse than the MIDAS model. The MIDAS model for U.S. crude oil inventories does not have systematically lower MSPE than the corresponding forecasting model based on monthly U.S. inventory data, however, and has comparable directional accuracy.

While the improvements in forecast accuracy are less substantial for other weekly financial predictors, the pattern of results is similar. Although MIDAS models often significantly outperform the no-change forecast, so do the corresponding forecasts from

models based on monthly financial predictors, and there is little to choose between these models. Examples include models based on oil futures prices, returns on oil stocks and gasoline price spreads. In some cases, the MIDAS model forecasts actually are inferior to the forecasts from the corresponding monthly model or they fail to improve on the no-change forecast.

Even when MIDAS models work well, therefore, not much is lost by ignoring high-frequency financial data in forecasting the monthly real price of oil. This finding is not only important for applied oil price forecasters, but also interesting from a methodological point of view. It reminds us that, despite the intuitive appeal of MIDAS models, it is by no means a foregone conclusion that the use of weekly predictors will improve the accuracy of monthly forecasts. The determining factor is whether the additional signal contained in the weekly data compensates for the additional noise. Different empirical applications can produce different results.

The remainder of the paper is organized as follows. In section 2 we review our data sources and the conventions used in transforming the daily data to weekly frequency. Section 3 provides a brief summary of the mixed-frequency forecasting models. Section 4 motivates the choice of the high-frequency predictors and contains the empirical results. Concluding remarks are provided in section 5.

2 Data

Our objective is to compare the real-time out-of-sample forecast accuracy for the monthly real price of oil of a set of models that include high-frequency financial market data. We focus on forecasts of the real U.S. refiners' acquisition cost of crude oil imports, which is a widely used proxy for the global price of oil (see Alquist et al. 2013). The refiners' acquisition cost measures what refiners actually pay for the crude oil they purchase. We deflate this price by the U.S. consumer price index for all urban consumers.

2.1 Data Construction

Throughout the paper, we focus on data measured at the weekly frequency, even if daily data are available, for three reasons. First, there is a potential trade-off between obtaining additional information and encountering noise in the high-frequency data. The use of weekly data strikes a balance in this regard. Second, in the early part of the sample there are gaps in the daily data for some of the time series that we consider. By relying on weekly data, we are able to construct internally consistent time series for longer time spans. Third, some of our data are available only at weekly frequency, and the choice of weekly data facilitates comparisons across forecasting models.

A complication that arises with weekly data is that some months consist of five instead of four weeks. We follow the approach proposed by Hamilton and Wu (2013) to generate a balanced weekly data set where each month consists of four weeks.² We use the observation of the last (trading) day of the week to convert daily data to weekly frequency. For the models estimated at monthly frequency, we take averages of daily data over the month, consistent with the construction of the U.S. Energy Information Administration (EIA) oil price data.

2.2 Data Sources

The daily West Texas Intermediate (WTI) spot oil price is obtained from the *Wall Street Journal* and the corresponding daily NYMEX oil futures prices for maturities of 1 to 18 months are obtained from Bloomberg.³ Daily data for the spot price of regular gasoline for delivery in New York Harbor are available from the EIA for the period June 1986

²For a Bayesian approach to modelling irregularly spaced data, see Chiu et al. (2012). It is unlikely that there would be gains from having one more weekly observation at irregular intervals in our models, because several alternative timing conventions that we considered generated very similar results.

³The spot price data start in January 1985, the oil futures price data for maturities 1 through 9 months start in June 1984, those for the 12-month maturity start in December 1988, those for the 15-month maturity start in June 1989 and those for the 18-month maturity start in October 1989.

to March 2013.⁴ The daily spot price index for non-oil industrial raw materials from the Commodity Research Bureau is available from June 1981 onwards. Daily data for the BDI are obtained from Bloomberg starting in January 1985. Data for U.S. crude oil inventories are reported from August 1982 onwards in the *Weekly Petroleum Status Report* issued by the EIA, but consistent weekly time series could be constructed only back to January 1984 due to gaps in the earlier data. Our analysis takes account of the fact that this report is issued every Wednesday and contains data extending to the preceding Friday. The closing price of the price-weighted NYSE Arca Oil Index is available from *Yahoo! Finance* from September 1983 onwards. This index is designed to measure the performance of the oil industry through changes in the stock prices of a cross-section of widely held corporations involved in the exploration, production and development of petroleum.⁵ Daily data for the closing price of the NYSE composite index, which measures the performance of all common stocks listed on the New York Stock Exchange, are obtained from *Yahoo! Finance* for the period January 1966 to March 2013. Weekly data for the federal funds rate, the 3-month LIBOR rate and the nominal trade-weighted U.S.-dollar index for major currencies are available from the Federal Reserve Economic Data (FRED) database from, respectively, July 1954 onwards, January 1986 onwards and January 1973 onwards.

The monthly real-time data for world oil production, the Kilian (2009) index of global real economic activity, the nominal refiners' acquisition cost of imported crude oil, the U.S. consumer price index for all urban consumers and the proxy for global crude oil inventories are taken from the real-time database developed by Baumeister and Kilian (2012), which contains vintages from January 1991 to March 2013.

⁴The gasoline spot price is reported in U.S. dollars per gallon and is converted to U.S. dollars per barrel by multiplying the price by 42 gallons/barrel to make it compatible with the crude oil price (see Baumeister et al. 2013).

⁵The index is composed of the following companies: Anadarko Petroleum, BP plc, ConocoPhillips, Chevron, Hess, Marathon Oil, Occidental Petroleum, Petr, Phillips 66, Total SA, Valero Energy, and Exxon Mobil.

3 Mixed-Frequency Real-Time Forecasting Models

In this section we review the forecasting models that will be considered in section 4. The objective is to forecast the monthly real price of oil using weekly predictors. For expository purposes, it is useful to focus on mixed-frequency VAR (MF-VAR) models first, before discussing MIDAS models.

3.1 MF-VAR Forecasts

There are two approaches to estimating the MF-VAR model. One is to estimate the model in state-space representation (see, e.g., Schorfheide and Song 2012). The other approach is to stack the weekly predictors in a vector depending on the timing of their release (see Ghysels 2012). The main difference compared with the state-space representation is that there are no missing observations, since the model is estimated at monthly frequency, and standard estimation methods can be used. We focus on the latter approach.

3.1.1 MF-VAR Model Represented as a Stacked-Vector System

Denote by x_t^1 , x_t^2 , x_t^3 and x_t^4 the releases of the weekly variables in the first, second, third and fourth week of each month t . Define $z_t = [x_t^{w'}, x_t^{m'}]'$ where $x_t^w = [x_t^{1'}, x_t^{2'}, x_t^{3'}, x_t^{4'}]'$ and x_t^m is the vector of monthly variables including the log of the real price of oil. Then the variables in the system evolve according to the monthly VAR model

$$A(L)z_t = u_t, \quad (1)$$

where u_t is white noise and $A(L)$ denotes the autoregressive lag order polynomial. The model in equation (1) can be estimated by least-squares methods, as in the case of a single-frequency VAR model. Forecasts of the real price of oil at monthly horizons $h = 1, \dots, 24$ can be generated by iterating the recursively estimated VAR model forward conditional on the date t information set and converting the forecast of the monthly real price of oil from log-levels to levels.

3.2 Univariate Mixed-Frequency Forecasts

A more parsimonious approach to dealing with mixed-frequency data involves specifying a univariate MIDAS regression. There are three alternative MIDAS representations. Let X_t^w denote a predictor observed in week $w \in \{1, 2, 3, 4\}$ of month t . The weekly predictor can depend on the horizon h of the forecast, in which case we add an additional superscript h . For example, we can define $X_t^{h,w}$ as the cumulative change in X_t^w between the last day of the current week and the last day of the same week h months ago. If the weekly predictor does not depend on h , the superscript h is dropped.

3.2.1 MIDAS Regression with Estimated Weights

The MIDAS model for combining weekly financial predictors with monthly oil price observations is defined as

$$R_{t+h} = R_t \left(1 + \beta B(L^{1/w}; \theta) X_t^w \right) + \varepsilon_{t+h}, \quad (2)$$

where R_t is the current level of the monthly real price of oil. The MIDAS lag polynomial $B(L^{1/w}; \theta)$ is an exponential Almon lag weight function

$$B(L^{1/w}, \theta) = \sum_{j=1}^4 b(j; \theta) L^{(j-1)/w},$$

where the lag operator is defined as

$$L^{(j-1)/w}(X_t^w) = X_{t-(j-1)/w}^w,$$

and $\theta \equiv \{\theta_1, \theta_2\}$ such that

$$b(j; \theta) = \frac{\exp(\theta_1 j + \theta_2 j^2)}{\sum_{j=1}^4 \exp(\theta_1 j + \theta_2 j^2)}.$$

Our results are not sensitive to the choice of the exponential Almon lag polynomial. Similar results would be obtained with a beta lag polynomial. The model parameters β and θ are recursively estimated by the method of non-linear least squares, and forecasts are generated as

$$R_{t+h|t} = R_t \left(1 + \hat{\beta} B(L^{1/w}; \hat{\theta}) X_t^w \right).$$

In some cases, there will be a priori reasons to restrict β to unity, in which case only θ has to be estimated.⁶

3.2.2 Equal-Weighted MIDAS Regressions

An even more parsimonious representation imposes equal weights on the weekly data, resulting in the MIDAS model:

$$R_{t+h} = R_t \left(1 + \beta \sum_{i=0}^3 \frac{1}{4} X_{t-i/4}^w \right) + \varepsilon_{t+h}. \quad (3)$$

In this case, no estimation is required except for the parameter β . The model is linear in β and can be estimated by ordinary least squares. If β is known, no regression is required and the MSPE of this model can be evaluated using the Diebold and Mariano (1995) test (the DM test).

3.2.3 Unrestricted MIDAS Regressions

Whether the added parsimony of the equal-weighted MIDAS model reduces the MSPE is an empirical question. An alternative approach is to relax the restrictions implied by the original MIDAS model. This yields the unrestricted MIDAS (or U-MIDAS) model:

$$R_{t+h} = R_t \left(1 + \sum_{i=0}^3 \alpha_i X_{t-i/4}^w \right) + \varepsilon_{t+h}. \quad (4)$$

Model (4) is linear in α_i and can be estimated by ordinary least squares.

4 Empirical Results

All forecasts are constructed subject to real-time data constraints. Unknown model parameters are estimated recursively. The forecast evaluation period starts in January 1992 and ends in September 2012. The use of such a long evaluation period minimizes the

⁶Note that the MIDAS model does not include an intercept. This fact allows us to nest the random walk forecast without drift. It can be shown that the inclusion of an intercept would systematically lower the forecast accuracy of our MIDAS models.

dangers of spurious forecast successes. The real oil price forecasts are evaluated in levels against the value of the real price of oil realized in the March 2013 vintage of the real-time data set. We discard the last six observations of the oil price data, which are still subject to revisions.

All forecasts are evaluated based on their MSPE relative to the MSPE of the monthly no-change forecast of the level of the real price of oil. MSPE ratios below 1 indicate that the model in question is more accurate than the no-change forecast. We also report the directional accuracy of the forecasts in the form of the success ratio, defined as the proportion of times that the model in question correctly predicts whether the real price of oil rises or falls. Under the null hypothesis of no directional accuracy one would expect a success ratio of 0.5. Higher ratios indicate an improvement on the no-change forecast.

While there is no valid test for the statistical significance of the real-time MSPE reductions from models based on estimated MIDAS or U-MIDAS weights, the equal-weighted MIDAS specification with $\beta = 1$ imposed does not suffer from parameter estimation uncertainty, allowing the use of conventional DM tests of equal MSPEs (see Diebold and Mariano 1995).⁷ The statistical significance of gains in directional accuracy is evaluated using the test of Pesaran and Timmermann (2009).

⁷The reason that we can assess the statistical significance of only the directional accuracy statistics and not the MSPE reductions is twofold. One problem is that all standard tests of equal MSPEs are based on the population MSPE, not the actual out-of-sample MSPE. This means that these tests are inappropriate for our purpose. This point was first made in Inoue and Kilian (2004) and has become widely accepted in recent years. If one uses these tests anyway, one will reject the null of equal MSPEs too often. This point has been illustrated, for example, in Alquist et al. (2013). Ongoing work by Clark and McCracken (2012) is addressing this issue, but their solutions do not apply in our context. The second problem is that standard tests for equal predictive accuracy do not apply when using real-time data. Clark and McCracken (2009) show how this problem can be overcome in the context of standard tests of no predictability in population. They focus on special cases under additional assumptions, but their analysis does not cover our forecast settings, nor does it address the first problem above.

4.1 MIDAS Results

The set of high-frequency predictors includes (1) the spread between the spot prices of gasoline and crude oil; (2) the spread between the oil futures price and the spot price of crude oil; (3) cumulative percent changes in the CRB index of the price of industrial raw materials, (4) U.S. crude oil inventories, and (5) the BDI; (6) returns and excess returns on oil stocks; (7) cumulative changes in U.S. nominal interest rates (LIBOR, federal funds rate); and (8) cumulative percent changes in the U.S. trade-weighted nominal exchange rate.

4.1.1 Oil Futures Prices

A good starting point are forecasting models based on oil futures prices. In the absence of a risk premium, arbitrage implies that the oil futures price is the conditional expectation of the spot price of oil (see Alquist and Kilian 2010). Equivalently, in logs this means that

$$E_t(\Delta s_{t+h}) = f_t^h - s_t, \quad (5)$$

where h denotes the forecast horizon and the maturity of the futures contract in months. For our sample period, the maximum maturity for which continuous time series of WTI oil futures and spot prices are available is 18 months. Expression (5) suggests that we define the MIDAS forecasting model for horizon h as a polynomial in $X_t^{h,w} = f_t^{h,w} - s_t^w$, where the spread is measured on the last day of week $w = 1, 2, 3, 4$ of a given month t . We also make an adjustment for expected inflation, which is approximated by the average inflation rate since July 1986, following Baumeister et al. (2013). As shown in Baumeister and Kilian (2012), this approximation is good enough in practice, given that the variation in the nominal price of crude oil far exceeds that in the inflation rate.

Table 1 shows that the equal-weighted MIDAS forecast has a lower MSPE than the no-change forecast at every horizon between 1 month and 18 months. The gains in accuracy are negligible at horizons under 12 months, but more substantial at longer horizons. The largest reduction in the MSPE is 17% at horizon 15. The MSPE reductions at horizons 12, 15 and 18 are statistically significant based on the DM test. There are no statistically significant gains in directional accuracy at short horizons. In fact, some of the success

ratios are well below 0.5. Significant improvements in directional accuracy are observed at horizons 9, 12, 15 and 18. The largest success ratio is 63%. Similar results are obtained for the model based on estimated MIDAS weights, and only slightly less accurate results for the unrestricted MIDAS model.

Although the MIDAS model compares favorably with the no-change forecast, so do traditional models based on the most recent monthly oil futures spread. The last two columns of Table 1 show the corresponding results based on the monthly oil futures model, as implemented in Baumeister and Kilian (2012). That model generates broadly similar results in that MSPE reductions are statistically significant at horizons 12 and 15, and directional accuracy at horizons 9, 12, 15 and 18. While the equal-weighted MIDAS model has a slightly lower MSPE at all horizons, the monthly forecasting model has a slightly higher and more statistically significant directional accuracy at longer horizons. Overall, there is little to choose between these models.

4.1.2 Gasoline Spreads

Petroleum products such as gasoline and heating oil are produced by refining crude oil. Many oil market analysts and financial analysts believe that the prices for these petroleum products contain useful information about the future evolution of the price of crude oil. In particular, changes in the product price spread – defined as the extent to which today's price of gasoline or heating oil deviates from today's price of crude oil – are widely viewed as a predictor of changes in the spot price of crude oil. For example, in April 2013 Goldman Sachs cut its oil price forecast citing significant downward pressure on product price spreads, which it interpreted as an indication of reduced final demand for products and hence an expectation of falling crude oil prices.

This forecasting approach has recently been formalized and evaluated by Baumeister, Kilian and Zhou (2013) using monthly data. Their analysis demonstrates that models of the gasoline price spread with an intercept of zero but a freely estimated slope parameter are reasonably successful at predicting the real price of oil at horizons up to 24 months. In the analysis below we impose the same restrictions. Preliminary analysis with alternative models confirmed that all other specifications are inferior.

Table 2 considers the MIDAS analogue of the model proposed in Baumeister et al. (2013), with $X_t^{h,w}$ denoting the spread between the spot price of gasoline and the WTI spot price of crude oil, measured on the last day of week $w = 1, 2, 3, 4$ of a given month t . The parameter β is freely estimated. Table 2 shows that this equal-weighted MIDAS model has a lower MSPE than the no-change forecast at every horizon from 1 month to 24 months, but with few exceptions the MSPE reductions are modest. There are no statistically significant gains in directional accuracy. Similar results hold when estimating the MIDAS weights. The unrestricted MIDAS model is somewhat less accurate.

Because of the presence of parameter estimation uncertainty, it is not possible to assess properly the statistical significance of the MSPE reductions in Table 2, but we can compare these results against those obtained for the corresponding monthly model, building on Baumeister, Kilian and Zhou (2013). The latter model has slightly lower MSPE at eight of the nine horizons. Both models' directional accuracy is statistically insignificant and erratic. There is no reason to favor one of these models. As in the case of the oil futures, there are no clear advantages to the use of the MIDAS model.

4.1.3 CRB Index of the Spot Price of Industrial Raw Materials

There is a long tradition of modelling oil prices jointly with other industrial commodities (e.g., Barsky and Kilian 2002; Frankel 2008). The Commodity Research Bureau (CRB) provides a widely used index of the spot price of industrial raw materials excluding crude oil. Alquist et al. (2013) first made the case that cumulative percent changes in this CRB price index in the recent past contain useful predictive information about expected changes in the price of oil. The rationale for this forecast is that, often, fluctuations in industrial commodity prices are driven by persistent and hence predictable variations in global real economic activity. Several studies have elaborated on this insight and demonstrated that such models have statistically significant directional accuracy and yield statistically significant MSPE reductions for the real price of oil (see Baumeister and Kilian 2012, 2013a,b).

The CRB index is also available on a daily basis, which allows us to incorporate weekly observations for the cumulative percent change in this index into a MIDAS model. The MIDAS model is estimated with $\beta = 1$ imposed. Table 3 shows that the equal-

weighted MIDAS model has directional accuracy at all horizons and statistically significant directional accuracy at some horizons. This model also reduces the MSPE at short horizons by as much as 14%, but the reductions are never statistically significant based on the DM test. At longer horizons there are no reductions in the MSPE. Similar results are obtained for the MIDAS model with estimated weights. The unrestricted MIDAS model is somewhat less accurate.

The last entries in Table 3 allow us to compare the performance of the MIDAS model to that of the corresponding model based on the monthly CRB predictor. The MSPE results are very similar and again statistically insignificant, but overall the monthly model has somewhat higher and more statistically significant directional accuracy. We conclude that in this case there is no gain from switching to MIDAS models and the monthly model is preferred.

4.1.4 Baltic Dry Index

The central idea behind using the CRB spot price index for industrial raw materials in forecasting the price of oil is that the real price of oil is predictable to the extent that the global business cycle is predictable. This is also the motivation for the inclusion of measures of global real economic activity such as the Kilian (2009) index in VAR oil price forecasting models. One limitation of the latter index as well as all other measures of global real economic activity is that it is not available at daily frequency. While there are daily real-time indices of U.S. real economic activity such as the business cycle conditions index of Aruoba, Diebold and Scotti (2009), there are no similar indices with the same global coverage as the monthly Kilian (2009) index.

An alternative business cycle indicator widely used by practitioners is the Baltic Dry Index (BDI), which is quoted on a daily basis by Bloomberg. This index is available starting in 1985. The name of this index derives from the fact that it is maintained by the Baltic Exchange in London. The BDI measures the cost of moving bulk dry cargo on representative ocean shipping routes in the world. Because dry bulk cargo primarily consists of materials that serve as industrial raw materials such as coal, steel, cement, and iron ore, this index is seen in the business world as an indicator of future industrial

production. In short, the BDI is viewed as a real-time leading economic indicator for the world economy and is used to predict future economic activity (e.g., Bakshi, Panayotov and Skoulakis 2011). This fact also makes it a potentially useful predictor of the real price of oil.

Despite its popularity among practitioners, the BDI differs in several dimensions from other measures of real economic activity based on dry cargo shipping rates, such as the Kilian (2009) index. Without further transformations, the BDI is at best a crude proxy for changes in global real economic activity. For the purpose of exploring its predictive content within the MIDAS framework, we focus on the percent change in the BDI over the last h months, rather than transforming the BDI into a business cycle index. The β parameter is freely estimated.

Table 4 shows that there is little gain in accuracy from including the BDI data. Apart from a negligible reduction in their MSPEs at the 1-month horizon, the first two MIDAS models tend have higher MSPEs than the random walk and lack directional accuracy at all horizons. The unrestricted MIDAS model is even less accurate. We conclude that there does not appear to be useful predictive information in the BDI data. This result is confirmed by the corresponding monthly regression models. Our findings underscore the importance of transforming the raw BDI data into a measure of global real activity prior to constructing oil price forecasts.

4.1.5 U.S. Crude Oil Inventories

Economic theory suggests that changes in expectations about the real price of oil all else equal are reflected in changes in crude oil inventories (see Alquist and Kilian 2010). This line of reasoning has led to the development of structural oil market models that explicitly model changes in global crude oil inventories (see Kilian and Murphy 2013; Kilian and Lee 2013; Knittel and Pindyck 2013). Monthly changes in global crude oil inventories also have been shown to have predictive power for the real price of oil (see Alquist et al. 2013). Although such data are not available at weekly frequency, U.S. crude oil inventories are. This fact suggests that we include percent changes in weekly U.S. crude oil inventories over the most recent h months in a MIDAS forecasting model for the real price of oil.

Table 5 considers two classes of MIDAS regressions. In the upper panel, we estimate the β parameter of the MIDAS model, whereas in the lower panel we impose $\beta = 1$ in estimation. This restriction improves the directional accuracy of the MIDAS model at longer horizons, while increasing the MSPE. Broadly speaking, the equal-weighted MIDAS model with $\beta = 1$ imposed is not much more accurate than the no-change forecast at short horizons, but substantially more accurate at longer horizons. Both the MSPE reductions and the improvements in directional accuracy are highly statistically significant.

This result is important because it suggests that the even-larger MSPE reductions for the model with β estimated are also likely to be statistically significant. The MIDAS model based on equal weights with β freely estimated is essentially tied with the no-change forecast at horizons 1, 3 and 6, but at higher horizons reduces the MSPE by up to 28% compared with the no-change forecast. The corresponding MIDAS model with $\beta = 1$ generates MSPE reductions only as high as 16%, but has higher and more statistically significant directional accuracy, making it the preferred model overall. Very similar, but marginally more accurate, results are obtained when the MIDAS weights are estimated. The unrestricted MIDAS model also performs well.

All MIDAS models have high and statistically significant directional accuracy, especially at longer horizons. The directional accuracy can be as high as 80%, which means that in 4 of 5 cases the model correctly predicts whether the real price of oil will go up or down. These estimates are higher than in any previous empirical study of oil price forecasting. We conclude that the equal-weighted MIDAS models based on weekly U.S. oil inventories are particularly promising tools for applied oil price forecasters.

Compared with the corresponding models based on monthly U.S. inventory data, however, the conclusion is less clear.⁸ Table 5 shows that the MIDAS model with $\beta = 1$ imposed tends to have a lower MSPE at all horizons, but only slightly so, whereas the MIDAS model with β estimated has a slightly higher or slightly lower MSPE than the monthly model, depending on the horizon. Likewise, there is little to choose between the monthly model and the MIDAS model when it comes to directional accuracy. Both models

⁸The monthly forecasting models are recursively estimated over the same estimation period as the MIDAS models.

perform quite well, especially at longer horizons.

4.1.6 Oil-Company Stock Prices

Chen (2013) recently showed that oil-sensitive stock price indices, particularly stock prices of oil companies, help forecast the real price of crude oil at short horizons. Such information is readily available at daily frequency. Building on Chen (2013), we explore this insight using a MIDAS regression with X_t^w denoting the weekly return on the NYSE Arca Oil Index, measured on the last day of week $w = 1, 2, 3, 4$ of a given month t . This index includes 13 major international oil and natural gas companies. The parameter β is freely estimated.

The upper panel of Table 6 shows that the MIDAS model with equal weights systematically reduces the MSPE relative to the no-change forecast for horizons up to 15 months. The largest MSPE reduction is 6% at the 1-month horizon. There also is some evidence of directional accuracy, but only the 1-month-ahead success ratio is statistically significant. When estimating the weights and when estimating the MIDAS model in its unrestricted form, however, the MSPE ratios deteriorate. Although the MIDAS model with equal weights performs better than the no-change forecast, it is not systematically more accurate than the monthly real-time forecast.⁹ There is no reason to prefer one specification over the other.

The lower panel of Table 6 shows that the same ranking of models applies when defining X_t^w as the weekly excess return on the NYSE Arca Oil Index relative to the NYSE Composite Index, except that the reductions in the MSPE and the improvements in directional accuracy are negligible.

⁹These reductions in the MSPE are considerably lower than those reported in Chen (2013). For example, Chen reports a 22% MSPE reduction at the 1-month horizon. These results differ for a number of reasons. First and most importantly, we are forecasting the real U.S. refiners' acquisition cost for crude oil imports, which is subject to real-time delays and revisions, whereas Chen (2013) focuses on the real WTI price, which for the most part is not subject to such delays and revisions. This accounts for about two-thirds of the difference in results. The remainder is largely accounted for by the fact that we focus on the monthly average price, as reported by the U.S. Energy Information Administration, rather than the end-of-month price that Chen focuses on.

4.1.7 U.S. Interest Rates

There is a perception among many observers that lower interest rates are associated with looser economic policies and hence higher demand for crude oil and possibly a lower supply of crude oil. Either way, this argument suggests a predictive relationship between changes in interest rates and changes in the price of oil. This perception has been boosted by studies suggesting that low real interest rates lead to high real commodity prices (see, e.g., Barsky and Kilian 2002; Frankel 2008).¹⁰ We investigate this proposition by fitting a MIDAS model for the difference between the interest rate on the last day of the current week and the interest rate h months earlier. We consider two alternative measures of U.S. interest rates. One is the U.S. federal funds rate, the other is the LIBOR rate. The parameter β is freely estimated.

Table 7 indicates that the approach yields modest MSPE reductions at horizons of 6 to 18 months for all MIDAS specifications involving the federal funds rate, but typically lacks directional accuracy. The corresponding results for the LIBOR rate are even less favorable, regardless of the specification. A comparison with the corresponding monthly forecasting model shows that very similar or worse results are obtained using monthly data only. Neither forecasting approach appears superior to the no-change forecast. This evidence reinforces skepticism regarding the empirical content of models linking oil price fluctuations to variations in U.S. interest rates. While there is no doubt about the theoretical link in question, its quantitative importance has yet to be established.

4.1.8 Trade-Weighted U.S. Exchange Rate

Another popular view is that fluctuations in the value of the dollar relative to other currencies predict changes in the real price of oil, as it becomes more or less expensive for importers of crude oil abroad to purchase crude oil. Previous studies of this question have found no evidence in monthly data to support this view (see Alquist et al. 2013).

Here we return to this question using MIDAS regression specifications that allow the use

¹⁰This argument is distinct from the implications of the Hotelling (1931) model of exhaustible resources that the price of oil should grow at the rate of interest. The latter proposition was evaluated and rejected in Alquist et al. (2013).

of high-frequency measures of the trade-weighted U.S. nominal exchange rate.

Table 8 shows that none of the MIDAS models produce reductions in the MSPE, although there is some evidence of directional accuracy at selected horizons. Exactly the same pattern applies to the corresponding monthly model in Table 8. There is some evidence of modest statistically significant directional accuracy at intermediate horizons, but again the MIDAS model has no advantage over the monthly model. We conclude that these models are indistinguishable.

Moreover, neither model can be recommended for forecasting oil prices, especially compared with some of the models discussed earlier. This result reinforces the conclusions in Alquist et al. (2013) about the lack of predictive content of exchange rates for oil prices. The notion that fluctuations in the trade-weighted U.S. exchange rate lead fluctuations in the real price of oil lacks empirical support.

4.2 MF-VAR Results

Despite the availability of numerous high-frequency predictors of the real price of oil, we conclude that only the weekly data on U.S. crude oil inventories stand out as useful predictors of the real price of oil. The surprisingly good performance of the MIDAS model based on U.S. crude oil inventories raises the question of whether even more accurate real-time forecasts could be obtained by incorporating the same weekly inventory data into an MF-VAR model.

Our baseline VAR model includes the percent change in global crude oil production, a measure of the global real activity proposed in Kilian (2009), the real price of oil and the change in global crude oil inventories. This choice of variables is motivated by economic theory (see Kilian and Murphy 2013; Kilian and Lee 2013). The model specification is identical to the specification employed in Baumeister and Kilian (2012), except that the lag order is restricted to 2 lags compared to 12 lags in the original analysis, since the MF-VAR model becomes computationally intractable for higher lag orders. By construction, in the MF-VAR(2) model there will be two months' worth of lags of the weekly predictor.

The results reported in Table 9 are based on the stacked vector representation of the mixed-frequency VAR model. Estimating the state-space representation of the model as

in Schorfheide and Song (2012) yields similar results (which are not shown, to conserve space). Table 9 shows that including weekly U.S. crude oil inventory data in the VAR(2) model does not improve the accuracy of the real-time VAR forecast. In fact, the MF-VAR(2) forecast is slightly less accurate than the original VAR(2) forecast. Either way, the MSPE reductions relative to the no-change forecast are small and do not extend beyond the 1-month horizon.

This evidence may seem to suggest that the information conveyed by the U.S. inventory data is already contained in the baseline VAR because of the inclusion of monthly global crude oil inventories. However, the corresponding MIDAS model in Table 5, which does not contain information about global crude oil inventories, is much more accurate than the VAR(2) model, especially at longer horizons, which suggests that the more parsimonious MIDAS model structure is what makes the difference. In fact, regardless of which high-frequency predictor is included in the MF-VAR(2) model, the MF-VAR(2) forecast rarely outperforms the random walk even at horizon 1 and never beyond horizon 3.¹¹ Our results indicate that MF-VAR models are systematically less accurate than MIDAS models in forecasting the real price of oil in real time.

5 Conclusion

We conclude that the best way to model mixed-frequency data in our context involves the use of MIDAS models rather than MF-VAR models. In general, the equal-weighted MIDAS model and the MIDAS model with estimated weights generate the most accurate real-time forecasts based on mixed-frequency data. We found no evidence that unrestricted MIDAS model forecasts are as accurate as, or more accurate than, forecasts from other MIDAS specifications.

Based on these MIDAS models, we reviewed a wide range of potential high-frequency financial predictors of the real price of oil. The results can be classified as follows:

- In many cases, the equal-weighted MIDAS model forecasts improve on the no-change forecast, but so does the corresponding forecast from a model including only lagged

¹¹These results are not shown to conserve space.

monthly data, and there is little to choose between the MIDAS model forecast and the forecast from the monthly model. Examples include models incorporating weekly oil futures spreads, weekly gasoline product spreads, weekly returns on oil company stocks and weekly U.S. crude oil inventories.

- In some cases, the MIDAS forecast improves on the no-change forecast somewhat, but is in turn inferior to the corresponding monthly real-time forecast. An example is the model incorporating cumulative percent changes in the weekly CRB spot price index for non-oil industrial raw materials.

- In yet other cases, the MIDAS forecast is about as accurate as the corresponding monthly forecast, but neither is systematically more accurate than the no-change forecast. Examples include models based on cumulative percent changes in the trade-weighted nominal U.S. exchange rate, the U.S. interest rates or the BDI.

Although many MIDAS models improve on the no-change forecast, the only time we documented large, systematic and statistically significant improvements in forecast accuracy occurred when we included weekly data on U.S. crude oil inventories in the MIDAS model. The latter specification not only yielded impressive reductions in the MSPE at horizons between 12 and 24 months, but also unusually high directional accuracy. The largest reduction in the MSPE we observed was 28%, and the largest success ratio was 80%. These gains in real-time forecast accuracy are large compared with those reported in any previous study on forecasting oil prices.

While our analysis produced strong new evidence that the monthly real price of oil is predictable at horizons beyond one year, this success cannot be attributed to the use of the MIDAS model, because the corresponding forecasting model based on monthly U.S. crude oil inventory data produces similar gains in accuracy. Our analysis suggests that, in general, unlike in many other studies, not much will be lost by ignoring high-frequency financial data in forecasting the monthly real price of oil.

Throughout this paper, we focused on MIDAS models for one high-frequency predictor at a time. An alternative strategy would have been to impose a factor structure on the set of high-frequency financial predictors, as in Andreou, Ghysels and Kourtellis (2013). The latter approach is natural in the context of macroeconomic forecasting, but less appealing

in our context given the much smaller number of potential predictors that can be motivated on economic grounds. The reason is that the real price of oil is determined in global oil markets and the set of relevant predictors is much smaller.

There are a number of potential extensions of our analysis. For example, although we focused on monthly oil price forecasts, it would have been straightforward to extend our analysis to quarterly horizons. Baumeister and Kilian (2013a,b) show that, usually, the best way to generate quarterly forecasts is to average monthly forecasts by quarter. One could also extend the analysis to include other oil price measures such as the WTI price. Doing so would raise additional complications discussed in Baumeister and Kilian (2013a). We focused on the real U.S. refiners' acquisition cost for crude oil imports in this paper because that price is a widely used proxy for the global price of oil.

References

1. Alquist, R., and L. Kilian (2010), "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, 25, 539-573.
2. Alquist, R., L. Kilian, and R.J. Vigfusson (2013), "Forecasting the Price of Oil," in: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland, 427-507.
3. Andreou, E., E. Ghysels, and A. Kourtellis (2010), "Regression Models with Mixed Sampling Frequencies," *Journal of Econometrics*, 158, 2, 246-261.
4. Andreou, E., E. Ghysels, and A. Kourtellis (2011), "Forecasting with Mixed-Frequency Data," in: Clements, M.P., and D.F. Hendry (eds.), *Oxford Handbook of Economic Forecasting*, 225-245.
5. Andreou, E., E. Ghysels, and A. Kourtellis (2013), "Should Macroeconomic Forecasters Use Daily Financial Data and How?" *Journal of Business and Economic Statistics*, 31, 240-251.
6. Aruoba, S.B., F.X. Diebold, and C. Scotti (2009), "Real-Time Measurement of Business Conditions," *Journal of Business and Economic Statistics*, 27, 417-427.

7. Bakshi, G., G. Panayotov, and G. Skoulakis (2011), "The Baltic Dry Index as a Predictor of Global Stock Returns, Commodity Returns, and Global Economic Activity," mimeo, University of Maryland.
8. Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: Bernanke, B., and K. Rogoff (eds.), *NBER Macroeconomics Annual 2001*, 137-183.
9. Baumeister, C., and L. Kilian (2012), "Real-Time Forecasts of the Real Price of Oil," *Journal of Business and Economic Statistics*, 30, 326-336.
10. Baumeister, C., and L. Kilian (2013a), "What Central Bankers Need to Know about Forecasting Oil Prices," forthcoming: *International Economic Review*.
11. Baumeister, C., and L. Kilian (2013b), "Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach," mimeo, University of Michigan.
12. Baumeister, C., L. Kilian, and X. Zhou (2013), "Are Product Spreads Useful for Forecasting? An Empirical Evaluation of the Verleger Hypothesis," mimeo, University of Michigan.
13. Bernard, J.-T., L. Khalaf, M. Kichian, and C. Yelou (2013), "On the Long-Term Dynamics of Oil Prices: Learning from Combination Forecasts," mimeo, Carleton University.
14. Chen, S.S. (2013), "Forecasting Crude Oil Price Movements with Oil-Sensitive Stocks," forthcoming: *Economic Inquiry*.
15. Chiu, C.W., B. Eraker, A.T. Foerster, T.B. Kim, and H.D. Seoane (2012), "Estimating VARs Sampled at Mixed or Irregularly Spaced Frequencies: A Bayesian Approach," mimeo, Federal Reserve Bank of Kansas City.
16. Clark, T.E., and M.W. McCracken (2009), "Tests of Equal Predictive Ability with Real-Time Data," *Journal of Business and Economic Statistics*, 27, 441-454.
17. Clark, T.E., and M.W. McCracken (2012), "Nested Forecast Model Comparisons: A New Approach to Testing Equal Accuracy," mimeo, Federal Reserve Bank of St. Louis.

18. Clements, M.P., and A.B. Galvao (2008), "Macroeconomic Forecasting with Mixed-Frequency Data: Forecasting US Output Growth," *Journal of Business and Economic Statistics*, 26, 546-554.
19. Clements, M.P., and A.B. Galvao (2009), "Forecasting US Output Growth Using Leading Indicators: An Appraisal Using MIDAS Models," *Journal of Applied Econometrics*, 24, 1187-1206.
20. Diebold, F.X., and R.S. Mariano (1995), "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics*, 13, 253-263.
21. Fattouh, B., L. Kilian, and L. Mahadeva (2013), "The Role of Speculation in Oil Markets: What Have We Learned So Far?" *Energy Journal*, 34, 7-33.
22. Foroni, C., E. Ghysels, and M. Marcellino (2013), "Mixed-Frequency Vector Autoregressive Models," forthcoming: *Advances in Econometrics*, 32.
23. Foroni, C., M. Marcellino, and C. Schuhmacher (2012), "U-MIDAS: MIDAS Regressions with Unrestricted Lag Polynomials," mimeo, EUI.
24. Frankel, J. (2008), "The Effect of Monetary Policy on Real Commodity Prices," in: Campbell, J. (ed.), *Asset Prices and Monetary Policy*, University of Chicago Press, 291-327.
25. Ghysels, E. (2012), "Macroeconomics and the Reality of Mixed Frequency Data," mimeo, University of North Carolina.
26. Ghysels, E., and J.H. Wright (2009), "Forecasting Professional Forecasters," *Journal of Business and Economic Statistics*, 27, 504-516.
27. Hamilton, J.D. (2008), "Daily Monetary Policy Shocks and New Home Sales," *Journal of Monetary Economics*, 55, 1171-1190.
28. Hamilton, J.D., and J.C. Wu (2013), "Risk Premia in Crude Oil Futures Prices," forthcoming: *Journal of International Money and Finance*.
29. Hotelling, H. (1931), "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39, 137-175.

30. Inoue, A., and L. Kilian (2004), "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?" *Econometric Reviews*, 23, 371-402.
31. Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053-1069.
32. Kilian, L. and T.K. Lee (2013), "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories," forthcoming: *Journal of International Money and Finance*.
33. Kilian, L., and D.P. Murphy (2013), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," forthcoming: *Journal of Applied Econometrics*.
34. Knittel, C.R., and R.S. Pindyck (2013), "The Simple Economics of Commodity Price Speculation," mimeo, MIT.
35. Pesaran, M.H., and A. Timmermann (2009), "Testing Dependence among Serially Correlated Multicategory Variables," *Journal of the American Statistical Association*, 104, 325-337.
36. Schorfheide, F., and D. Song (2012), "Real-Time Forecasting with a Mixed Frequency VAR," mimeo, University of Pennsylvania.

Table 1: Forecasting the monthly real price of oil with the oil futures spread
Evaluation period: 1992.1-2012.9

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.996	0.478	1.000	0.466	1.014	0.466	0.997	0.462
3	0.965	0.530	0.954	0.563	0.941	0.571	0.974	0.498
6	0.975	0.488	0.964	0.508	0.980	0.488	0.975	0.512
9	0.938	0.568*	0.922	0.564	0.939	0.568	0.944	0.589*
12	0.872*	0.592*	0.857	0.601*	0.878	0.601*	0.886**	0.613*
15	0.829*	0.621*	0.829	0.617*	0.890	0.638*	0.860**	0.634*
18	0.848*	0.629*	0.854	0.625*	0.962	0.625*	0.906	0.621*

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w} - E_t(\pi_t^h)))$, Equal weights
- $R_{t+h|t} = R_t(1 + B(L^{1/4}; \hat{\theta})(X_t^{h,w} - E_t(\pi_t^h)))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w} - E_t(\pi_t^h)))$, Unrestricted
- $R_{t+h|t} = R_t(1 + X_t^h - E_t(\pi_t^h))$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the difference between the log of the oil futures price for maturity h and the log of the spot price of oil in week w of month t , X_t^h is the difference between the log of the oil futures price for maturity h and the log of the spot price of oil in month t , and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using *(5% significance level) and **(10% significance level). For the equal-weighted MIDAS model and for the monthly model, statistically significant reductions in the MSPE according to the Diebold-Mariano test are marked using *(5% significance level) and **(10% significance level).

Table 2: Forecasting the monthly real price of oil with the gasoline-crude oil spot price spread
Evaluation period: 1992.1-2012.9

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.993	0.578	0.998	0.590	1.071	0.554	0.989	0.562
3	0.996	0.583	1.004	0.583	1.019	0.534	0.990	0.583
6	0.991	0.574	0.984	0.582	0.997	0.533	0.978	0.545
9	0.984	0.490	0.987	0.494	1.011	0.485	0.963	0.436
12	0.963	0.441	0.961	0.483	0.964	0.555	0.934	0.521
15	0.956	0.532	0.950	0.540	0.945	0.591	0.931	0.516
18	0.973	0.504	0.970	0.543	0.966	0.582**	0.971	0.470
21	0.976	0.541	0.972	0.563	1.003	0.546	0.986	0.454
24	0.935	0.588	0.927	0.566	0.953	0.540	0.934	0.500

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, Equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}) - E_t(\pi_t^h))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, Unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h - E_t(\pi_t^h))$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the difference between the log of the gasoline spot price and the log of the spot price of oil in week w of month t , X_t^h is the difference between the log of the gasoline spot price and the log of the spot price of oil in month t , and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using * (5% significance level) and ** (10% significance level).

Table 3: Forecasting the monthly real price of oil with the CRB spot price index of industrial raw materials
Evaluation period: 1992.1-2012.9

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.929	0.558**	0.927	0.562	0.978	0.546	0.934	0.546**
3	0.862	0.628*	0.831	0.636*	0.861	0.632*	0.863	0.628*
6	1.113	0.611*	1.112	0.623*	1.085	0.570*	1.107	0.598*
9	1.163	0.573	1.158	0.564	1.085	0.469	1.143	0.593*
12	1.132	0.546	1.131	0.546	1.131	0.454	1.100	0.592*
15	1.150	0.574**	1.144	0.574	1.131	0.451	1.118	0.617*
18	1.254	0.539	1.252	0.539	1.154	0.418	1.232	0.578*
21	1.382	0.528	1.382	0.528	1.139	0.445	1.376	0.528
24	1.377	0.513	1.380	0.509	1.172	0.451	1.394	0.443

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, Equal weights
- $R_{t+h|t} = R_t(1 + B(L^{1/4}; \hat{\theta})(X_t^{h,w}) - E_t(\pi_t^h))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, Unrestricted
- $R_{t+h|t} = R_t(1 + X_t^h - E_t(\pi_t^h))$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percent change in the CRB spot price index of industrial raw materials over the preceding h months in week w of month t , X_t^h is the percent change in the CRB spot price index of industrial raw materials over the preceding h months in month t , and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using * (5% significance level) and ** (10% significance level). For the equal-weighted MIDAS model and for the monthly model, statistically significant reductions in the MSPE according to the Diebold-Mariano test are marked using * (5% significance level) and ** (10% significance level).

Table 4: Forecasting the monthly real price of oil with the Baltic Dry Index
Evaluation period: 1992.1-2012.9

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.950	0.502	0.947	0.498	0.991	0.442	0.952	0.546**
3	1.049	0.470	1.079	0.466	1.078	0.482	1.109	0.462
6	1.015	0.504	1.023	0.512	1.083	0.520	1.056	0.492
9	1.030	0.502	1.025	0.498	1.033	0.490	1.124	0.548
12	1.087	0.445	1.094	0.445	1.166	0.441	1.447	0.500
15	1.123	0.383	1.136	0.387	1.203	0.391	1.544	0.426
18	1.297	0.435	1.308	0.414	1.327	0.397	2.112	0.474
21	1.399	0.341	1.393	0.332	1.397	0.358	2.214	0.411
24	1.391	0.363	1.407	0.363	1.464	0.442	2.185	0.327

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w}))$, Equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, Unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percent change in the BDI over the preceding h months in week w of month t , and X_t^h is the percent change in the BDI over the preceding h months in month t . The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using *(5% significance level) and **(10% significance level).

Table 5: Forecasting the monthly real price of oil with U.S. crude oil inventories
Evaluation period: 1992.1-2012.9

$\beta = \hat{\beta}$								
Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	1.000	0.530	0.998	0.550*	1.003	0.478	1.001	0.414
3	1.004	0.599*	1.008	0.587*	1.018	0.579*	0.998	0.575**
6	1.007	0.463	1.010	0.463	1.021	0.451	1.018	0.537
9	0.964	0.506	0.961	0.523	0.985	0.523	0.981	0.519
12	0.922	0.559	0.910	0.584	0.929	0.588	0.926	0.534
15	0.886	0.609**	0.881	0.613**	0.884	0.566	0.886	0.630**
18	0.835	0.621*	0.828	0.625*	0.836	0.599**	0.835	0.629**
21	0.688	0.712*	0.686	0.729*	0.690	0.734*	0.681	0.716*
24	0.720	0.686*	0.706	0.695*	0.714	0.712*	0.695	0.708*
$\beta = 1$								
Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.995	0.542	0.995	0.546**	1.003	0.478	1.001	0.586
3	0.985	0.599*	0.987	0.583*	1.018	0.579*	0.991	0.575**
6	1.013	0.512	1.016	0.500	1.021	0.451	1.025	0.537
9	0.945*	0.622*	0.943	0.622*	0.985	0.523	0.952*	0.610*
12	0.924*	0.710*	0.922	0.731*	0.929	0.588	0.930	0.744*
15	0.912*	0.766*	0.911	0.779*	0.884	0.566	0.916*	0.762*
18	0.899*	0.776*	0.898	0.780*	0.836	0.599**	0.904*	0.797*
21	0.848*	0.803*	0.848	0.830*	0.690	0.734*	0.847*	0.799*
24	0.844*	0.774*	0.842	0.783*	0.714	0.712*	0.844*	0.792*

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w}))$, Equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, Unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percent change in U.S. crude oil inventories over the preceding h months in week w of month t , and X_t^h is the percent change in U.S. crude oil inventories over the preceding h months in month t . The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using * (5% significance level) and ** (10% significance level). For the equal-weighted MIDAS model and for the monthly model in the lower panel, statistically significant reductions in the MSPE according to the Diebold-Mariano test are marked using * (5% significance level) and ** (10% significance level).

Table 6: Forecasting the monthly real price of oil with returns on oil stocks
Evaluation period: 1992.1-2012.9

Returns on the NYSE Oil Index								
MIDAS								
Horizon (months)	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.943	0.586*	0.987	0.570**	0.999	0.590	0.945	0.518
3	0.952	0.567	0.970	0.575**	0.972	0.567	0.951	0.547
6	0.986	0.529	0.991	0.545	0.998	0.537	0.984	0.504
9	0.986	0.523	1.000	0.531	1.022	0.560	0.989	0.531
12	0.986	0.576	1.004	0.571**	1.032	0.563	0.983	0.588*
15	0.991	0.515	0.999	0.528	1.024	0.536	0.990	0.506
18	1.004	0.496	1.008	0.435	1.026	0.453	1.018	0.483
21	1.003	0.476	1.015	0.463	1.017	0.463	1.007	0.459
24	0.994	0.447	1.007	0.509	0.995	0.496	1.002	0.465

Excess Returns of the NYSE Oil Index relative to the NYSE Composite Index								
MIDAS								
Horizon (months)	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.968	0.554*	1.007	0.538**	1.010	0.530**	0.973	0.530
3	0.982	0.518	0.998	0.518	1.001	0.522	0.985	0.526
6	0.993	0.496	0.999	0.537	1.003	0.520	0.996	0.508
9	1.002	0.469	1.023	0.502	1.033	0.535	1.002	0.486
12	1.000	0.500	1.019	0.534	1.046	0.521	0.998	0.517
15	0.999	0.485	1.011	0.532	1.048	0.489	1.001	0.502
18	1.004	0.478	1.015	0.483	1.037	0.427	1.001	0.500
21	1.000	0.502	1.015	0.441	1.026	0.450	0.997	0.520**
24	1.003	0.482	1.019	0.491	1.019	0.434	1.001	0.447

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w}))$, Equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, Unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the 1-week return (or excess return) on the NYSE oil index in week w of month t , and X_t^h is the 1-month return (or excess return) on the NYSE oil index in month t . The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using *(5% significance level) and **(10% significance level).

Table 7: Forecasting the monthly real price of oil with U.S. interest rates
Evaluation period: 1992.1-2012.9

Federal Funds Rate								
Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.998	0.510	0.999	0.534**	1.001	0.502	0.998	0.470
3	1.004	0.530	1.004	0.538	1.005	0.526	1.004	0.530
6	0.969	0.459	0.971	0.504	0.967	0.520	0.966	0.459
9	0.960	0.506	0.963	0.510	0.964	0.515	0.953	0.506
12	0.952	0.504	0.947	0.475	0.946	0.483	0.952	0.496
15	0.961	0.515	0.954	0.502	0.946	0.502	0.963	0.498
18	0.986	0.491	0.982	0.487	0.977	0.487	0.987	0.500
21	1.011	0.480	1.009	0.472	0.997	0.480	1.012	0.489
24	1.032	0.434	1.032	0.442	1.024	0.438	1.032	0.434

LIBOR								
Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	1.006	0.522	1.010	0.526	1.013	0.530	1.037	0.534**
3	1.017	0.538**	1.018	0.506	1.018	0.571*	1.023	0.547
6	0.996	0.463	0.996	0.475	1.033	0.496	1.014	0.385
9	0.994	0.436	0.992	0.461	0.992	0.461	1.086	0.486
12	0.980	0.458	0.979	0.454	0.986	0.483	1.050	0.382
15	0.995	0.485	0.994	0.481	0.993	0.485	1.033	0.430
18	1.011	0.457	1.011	0.461	1.008	0.470	1.050	0.457
21	1.033	0.459	1.034	0.454	1.034	0.480	1.083	0.389
24	1.058	0.429	1.060	0.434	1.064	0.460	1.088	0.358

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w}))$, Equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, Unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the change in the interest rate over the preceding h months in week w of month, and X_t^h is the change in the interest rate over the preceding h months in month t . The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using * (5% significance level) and ** (10% significance level).

Table 8: Forecasting the monthly real price of oil with the nominal trade-weighted U.S. exchange rate
Evaluation period: 1992.1-2012.9

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	1.005	0.466	1.006	0.514	1.018	0.514	1.007	0.466
3	1.081	0.502	1.078	0.486	1.084	0.490	1.068	0.494
6	1.006	0.426	1.016	0.418	1.038	0.434	1.000	0.480
9	1.061	0.622*	1.070	0.548	1.097	0.523	1.069	0.618*
12	1.174	0.618*	1.188	0.613*	1.199	0.592*	1.176	0.626*
15	1.149	0.591*	1.147	0.600*	1.176	0.600*	1.146	0.600**
18	1.157	0.565	1.163	0.547	1.175	0.543	1.153	0.560
21	1.143	0.459	1.146	0.472	1.163	0.472	1.140	0.463
24	1.079	0.482	1.079	0.451	1.078	0.465	1.078	0.478

NOTES: The forecasts are constructed as:

- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^3 \frac{1}{4}(X_{t-i/4}^{h,w})) + \epsilon_t$, Equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w})) + \epsilon_t$, Estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^3 \hat{\alpha}_i(X_{t-i/4}^{h,w})) + \epsilon_t$, Unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, Monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the change in the exchange rate over the preceding h months in week w of month t , and X_t^h is the change in the exchange rate over the preceding h months in month t . The benchmark model is the monthly no-change forecast. Boldface indicates improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using * (5% significance level) and ** (10% significance level).

Table 9: VAR and MF-VAR forecasts of the monthly real price of oil
Evaluation period: 1992.1-2012.9

Horizon (months)	VAR(2)		MF-VAR(2) with weekly U.S. crude oil inventories	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	0.915	0.566*	0.950	0.530
3	1.007	0.543	1.090	0.522
6	1.108	0.553	1.244	0.459
9	1.224	0.539	1.479	0.436
12	1.309	0.563	1.630	0.458
15	1.362	0.549	1.735	0.447
18	1.426	0.539	1.908	0.427
21	1.487	0.533	2.064	0.450
24	1.482	0.518	2.071	0.465

NOTES: The four variables in the VAR model are the growth rate of world oil production, the log of the real price of oil, the Kilian (2009) global real economic activity index, and the change in global crude oil inventories. The weekly U.S crude oil inventories are expressed as the percent change over the preceding h months. The benchmark model is the monthly no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran-Timmermann test are marked using *(5% significance level) and **(10% significance level).